

RESEARCH ARTICLE

Assessing the contagiousness of mass shootings with nonparametric Hawkes processes

Peter Boyd[☉], James Molyneux^{☉*}

Department of Statistics, Oregon State University, Corvallis, Oregon, United States of America

☉ These authors contributed equally to this work.

* molyneuj@oregonstate.edu

OPEN ACCESS

Citation: Boyd P, Molyneux J (2021) Assessing the contagiousness of mass shootings with nonparametric Hawkes processes. PLoS ONE 16(3): e0248437. <https://doi.org/10.1371/journal.pone.0248437>

Editor: Haroldo V. Ribeiro, Universidade Estadual de Maringá, BRAZIL

Received: October 8, 2020

Accepted: February 16, 2021

Published: March 30, 2021

Copyright: © 2021 Boyd, Molyneux. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: Our data comes for four different third-party sources are publicly available at GitHub here: https://github.com/boydpe/mass_shooting_contagion. Additionally, these data are within the paper's [Supporting information](#) file. The code used in creating our analysis, plots and tables can be found here: <https://github.com/boydpe/nphawkes>.

Funding: The author(s) received no specific funding for this work.

Competing interests: The authors have declared that no competing interests exist.

Abstract

Gun violence and mass shootings are high-profile epidemiological issues facing the United States with questions regarding their contagiousness gaining prevalence in news media. Through the use of nonparametric Hawkes processes, we examine the evidence for the existence of contagiousness within a catalog of mass shootings and highlight the broader benefits of using such nonparametric point process models in modeling the occurrence of such events.

Introduction

Gun violence in the United States is a national public health crisis [1] with firearm homicide rates 19.5 times that of other high-income countries [2]. Mass shootings in particular represent a phenomenon of interest in that these high-profile events with multiple, and occasionally numerous, victims generate large amounts of media coverage. Such media coverage may lead to both a contagion effect that may incite others to carry out similar acts as well as an imitation effect that may allow mass shooters to learn from those that preceded them [3]. Though the term *mass shooting* lacks a specific, rigorous definition, the number of gun related incidences with multiple victims has become so common in the past two decades that research of these events has become a necessary component of public health studies in the United States [4]. From 2000 to 2018, the US Federal Bureau of Investigation (FBI) recorded 277 active shooter incidents in which an individual shoots and kills (or attempts to kill) others in a public space, resulting in 2430 casualties [5]. The FBI further notes that the number of such incidents is on the rise, with 69% of these incidents occurring between 2010 and 2018. The need to address the contagion factor of these events, whereby a single mass shooting event inspires or is correlated with future mass shooting events, represents a fundamental question in the underlying mass shooting phenomenon.

Previous research proposed that the ideation and implementation of mass shootings are linked to media coverage of such events [6]. A contagion factor was also previously found by Towers, et al., [7] which used a self-excitation contagion model to quantify the degree to which previous events inspired future events. In their work, Towers, et al. model the increased probability of a mass shooting event occurring on day t_j given a previous event occurred on day t_i , t_i

$< t_j$, and the average duration of the contagion process T_{excite} using an exponential probability distribution. That is, the probability of a new mass shooting event occurring sometime in the 24 hours of day t_j is expressed as

$$P(t_j|t_i, T_{excite}) = \int_{t_j-t_i-1}^{t_j-t_i} dx \frac{e^{-x/T_{excite}}}{T_{excite}}. \quad (1)$$

Towers, et al. then couple this probability model with a non-contagion related baseline number of events, $N_0(t)$, and a total number of expected secondary events, $N_{secondary}$, to compute an expected number of events, N^{exp} , on day t_n expressed as

$$N^{exp}(t_n) = N_0(t_n) + N_{secondary} \sum_{i:t_i < t_n} P(t_n|t_i, T_{excite}). \quad (2)$$

By leveraging methods from the self-exciting point process literature, we propose to improve upon these previous studies in a few distinct but important ways. First, by formulating the occurrence of mass shootings as a nonparametric Hawkes process we avoid having to make assumptions about how the contagion factor decays over time. Whereas the decay of the temporary excitation in space, time, or space-time is well-understood in some fields, such as in seismology where the occurrence of aftershocks has been shown to decay according to a power-law, the decay of a potential contagion factor in mass shootings is less understood. Thus, by expressing the contagion factor nonparametrically we avoid having to assume that the decay of the contagion factor follows some prescribed probability distribution.

Secondly, by using the EM-type model independent stochastic declustering (MISD) method of Marsan and Lengliné [8], we can estimate the probability that any individual mass shooting event was caused by a previous event or is a non-contagion related background event. This in turn allows for the estimation of the background rate of mass shootings which can then be expressed as the expected number of background events by taking the sum of the background event probabilities. Further, the modeling framework allows for the expected number of secondary events to vary according to the heinousness of the crime as measured by the number of victims.

This article is then organized as follows. Both parametric and nonparametric self-exciting and Hawkes point processes will be introduced along with model-fit assessments. We then introduce the four mass shooting data sets which we utilize in the analyses. Finally, the significance of the results, comparisons to other analyses, and conclusions are discussed.

Methods

A point process is a random collection of points $\{\tau_1, \tau_2, \dots\}$ occurring in some metric space [9]. These points often occur in some temporal or spatio-temporal window where $t_i \in \mathbb{R}$ represents the temporal dimension of the i^{th} point and $s_i \in \mathbb{R}^n$ represents the spatial coordinates of the i^{th} point. In practice, \mathbb{R}^n is often taken to be \mathbb{R}^2 or \mathbb{R}^3 so that the spatial coordinates land on some two-dimensional plane or three-dimensional space where the third dimension can be taken to be the depth of the point. For our purposes, we consider the occurrence of mass shooting events to be a collection of n marked spatio-temporal points, $\{(t_i, x_i, y_i, m_i): i = 1, 2, \dots, n\}$, such that $t_i \in [0, T]$ represents the time the event occurred with 0 and T taken to be the start and end of the temporal window, respectively, and $(x_i, y_i) \in (-\infty, \infty) \times (-\infty, \infty)$ represents the spatial location of the event. The mark, m_i , of point i is then some additional covariate information which we take to be the number of victims, excluding the perpetrator, of the i^{th} mass shooting event. For the marks of the process, we define the number of victims to be the number of individuals either killed or injured during the shooting. In defining the marks in

this way, we intend to measure how events with different numbers of victims impacts the ability of an event to incite future events.

In general, point processes are typically modeled via their conditional intensity function, $\lambda(t)$ or $\lambda(s, t)$ for time and space-time point processes, respectively. The conditional intensity is defined as the infinitesimal expected rate at which points occur given the history of the processes, \mathcal{H}_t . That is, we model the occurrence of points in time as

$$\lambda(t|\mathcal{H}_t) = \lim_{\Delta t \rightarrow 0} \frac{E[N([t, t + \Delta t))|\mathcal{H}_t]}{\Delta t} \quad (3)$$

or in space-time as

$$\lambda(s, t|\mathcal{H}_t) = \lim_{\Delta s, \Delta t \rightarrow 0} \frac{E[N((s, s + \Delta s) \times (t, t + \Delta t))|\mathcal{H}_t]}{\Delta s \Delta t} \quad (4)$$

where $N(\cdot)$ is taken to be a counting measure [9].

In what follows, we introduce the self-exciting, or Hawkes, point process and then elaborate further on the estimation and evaluation of the nonparametric version of the processes.

Hawkes and self-exciting point processes

When the occurrence of a point causes the temporary elevation in the occurrence of future points nearby in time or space and time, we refer to such a process as a self-exciting point process. Foundational work in self-exciting point processes was performed by Hawkes [10] who defined the conditional intensity as

$$\lambda(t|\mathcal{H}_t) = \mu + \sum_{i:t_i < t} v(t - t_i). \quad (5)$$

where μ specifies the background rate in which events stochastically occur in time and v is the triggering function which governs the temporary self-excitation of events. A self-exciting point process can be categorized as a branching process, or a mathematical process in which a background event occurs and spawns additional offspring events, which can in turn have additional offspring of their own.

The temporal Hawkes process was later extended to the spatio-temporal domain where the rate of events can be modeled not just at time t but also location s . When considering spatio-temporal self-exciting point processes specifically, the conditional intensity function is defined as

$$\lambda(s, t|\mathcal{H}_t) = \mu(s) + \sum_{i:t_i < t} v(s - s_i, t - t_i). \quad (6)$$

where $\mu(s)$ describes the background rate for the occurrence of points as a function of the spatial location and $v(s - s_i, t - t_i)$ again describes the excitation of the events.

There are numerous applications for self-exciting point processes, most notably in seismology [11, 12], social networks such as email chains [13] or retweets on Twitter [14], criminology and gang related violence [15], terrorism [16], neuroscience [17], and the spread of epidemic diseases like Ebola [18]. In this paper, we focus specifically on the realization of self-exciting point processes as applied to mass shootings.

Epidemic-type aftershock sequences. Modeling the temporal occurrence of earthquakes with self-exciting point processes was first proposed by Ogata [11] who later extended his model to the spatio-temporal domain [12]. These epidemic-type aftershock sequence (ETAS) models are parametric models based on well-studied phenomenon of the temporal decay of

aftershocks, via Omori-Utsu [19, 20], and the magnitude distributions of earthquakes, via Gutenberg-Richter [21].

ETAS models consider the triggering function ν to be composed of three separable functions $g(t)$, $h(x, y)$, and $k(m)$ pertaining to time, space, and magnitude, respectively. That is, the conditional intensity function is defined as

$$\lambda(x, y, t | \mathcal{H}_t) = \mu(x, y) + \sum_{i: t_i < t} g(t - t_i) h(x - x_i, y - y_i) k(m_i) \quad (7)$$

The functions $g(t)$ and $h(x, y)$ then model how the conditional rate of events decays over time and space, respectively, while the function $k(m)$ describes the productivity of previous events based on their marks. In seismology, for example, where marks are taken to be the magnitude, or amount of energy released during an earthquake, events with a larger magnitude will be more productive at producing offspring than an event of smaller magnitude.

Nonparametric Hawkes models

Whereas ETAS models are based on well-understood properties of seismic phenomenon to parameterize the components of the triggering functions, mass shootings are much less well studied and understood. For this reason, we use a nonparametric Hawkes model [8] to study the occurrence and contagiousness of mass shootings. This modeling framework allows for a causal structure to be calculated probabilistically using an iterative process to estimate the probability that an event was caused by a previous event, or conversely is a background event. These probabilities are then used to estimate the constant values of step-functions for each portion of the triggering function, also known as histogram estimators. For the histogram estimators, the differences in the pairwise times and locations of the events are computed, and then each inter-event time and location is placed into a set of disjoint intervals or bins. Based on the probabilities of the iterative process, constant values are estimated for each interval in order to fit the model. The histogram estimator for the marks works similarly except that the disjoint intervals are created based on the marks themselves and not the pairwise differences.

Using a nonparametric Hawkes model to describe mass shootings then allows us to model the contagion factor of these events without making parametric assumptions about the shape and rate of decay. It also allows us to obtain estimates for the background rate of events directly based on the estimated probabilities of the model.

Following the work of [8], we define

$$p_{ij} = \begin{cases} \text{probability event } i \text{ is triggered by event } j, & i > j \\ \text{probability event } i \text{ is a background event,} & i = j \\ 0, & i < j. \end{cases} \quad (8)$$

These probabilities can then be displayed as a lower-triangular probability matrix P with

$$P = \begin{bmatrix} p_{11} & 0 & 0 & \cdots & 0 \\ p_{21} & p_{22} & 0 & \cdots & 0 \\ p_{31} & p_{32} & p_{33} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & 0 \\ p_{n1} & p_{n2} & p_{n3} & \cdots & p_{nn} \end{bmatrix}. \quad (9)$$

Each row, i , of the probability matrix then describes the probability that event i was caused by event j , $i > j$, or is itself a background event, $i = j$. Thus, each row of the probability matrix must sum to one.

After initializing the probability matrix P by setting each $p_{i\cdot} = 1/i$, we iterate over the following sequence of steps until convergence has been achieved:

1. Update the stationary background rate of the process by computing the expected number of background events based on the estimated probabilities from the P matrix.
2. Update the histogram estimators of the triggering functions for each disjoint space, time, or mark interval using the estimated probabilities from the P matrix.
3. Use the now updated background rate and triggering functions to update the probabilities that each event was either a background event or a child of a previous event.

Convergence is achieved once the largest update to the entries of the probability matrix falls below some prescribed value ε . For a more detailed description of the algorithm, we refer readers to the article by Fox, et al. [13].

Standard errors for the histogram estimators of the triggering functions can also be computed to assess the variability of the estimates. For the case of the temporal triggering function, $g(t)$, let S_ℓ denote a binomial random variable that represents the number of offspring in bin ℓ , defined by parameters η_t , the true number of offspring, and θ_ℓ^g , the true probability a triggered event falls in bin ℓ . We attain estimates of these parameters via

$$\hat{\eta}_t = \sum_{i=1}^n \sum_{j=1}^{i-1} p_{ij} \quad \text{and} \quad \hat{\theta}_\ell^g = \sum_{B_\ell} p_{ij} / \hat{\eta}_t \quad (10)$$

for p_{ij} equal to the triggering probability of the matrix P upon convergence of the MISD algorithm and B_ℓ equal to the set of all events whose time differences fall within bin ℓ [13]. As a result, we estimate the variance of the value $\hat{g}(t) = g_\ell = S_\ell / (\Delta t_\ell \eta_t)$ by

$$\widehat{Var}(g_\ell) = \frac{(\hat{\theta}_\ell^g)(1 - \hat{\theta}_\ell^g)}{\hat{\eta}_t \Delta t_\ell^2}. \quad (11)$$

Standard errors for $\hat{k}(m) = k_\ell$ can be computed similarly via

$$\widehat{Var}(k_\ell) = \frac{\hat{n}_t (\hat{\theta}_\ell^k)(1 - \hat{\theta}_\ell^k)}{(n_\ell^{mark})^2} \quad (12)$$

where $\hat{\theta}_\ell^k = \sum_{A_\ell} p_{ij} / \hat{\eta}_t$ for A_ℓ equal to the set of events whose marks fall within the ℓ^{th} marks bin and N_ℓ^{mark} equal to the number of events in bin ℓ .

Model evaluation via super-thinning

Super-thinning is a hybrid approach of two combined model evaluation techniques for point processes: residual thinning and superpositioning [22]. For residual thinning, event i is kept in the realized set of points, \mathcal{S} , with probability $b / \hat{\lambda}(s_i, t_i)$ for $b = \inf_{(s,t) \in \mathcal{S}} \{ \hat{\lambda}(s, t) \}$ and removed from the data otherwise, where $\hat{\lambda}$ represents the estimated conditional intensity of a point in \mathcal{S} [23]. Superpositioning meanwhile first simulates a point process with intensity $b - \hat{\lambda}(s, t)$,

where $b = \sup_{(s,t) \in \mathcal{S}} \{ \hat{\lambda}(s, t) \}$ and then superposes these points into the data [24]. For both

thinning and superpositioning, the resulting residual process, \mathcal{R} , will be a homogeneous Poisson process if and only if the model for the conditional intensity, λ , is correct. By using the hybrid approach of super-thinning, in that a point process is both thinned in areas of high conditional intensity and superposed points are included in areas of low intensity to form the residual process, the resulting set of points will have a higher power and lower volatility.

The process then for super-thinning a point process, \mathcal{S} , is as follows:

1. Thin \mathcal{S} by retaining each point (s_i, t_i) with probability $\min\{b/\hat{\lambda}(s_i, t_i), 1\}$.
2. Simulate a point process with rate $\max\{b - \hat{\lambda}(s, t), 0\}$ at point (s, t) .
3. Combine the two resulting processes above to form the super-thinned residual process, \mathcal{R} .

The value of b is used to adjust how much thinning or superposing takes place. Once a residual process is obtained, it can be examined for uniformity. If the model specification is correct, then the residual process should have a uniform distribution throughout the time window.

Mass shooting data

Data availability on mass shootings is limited with no definitive collection of incidents reported by a public entity, in part due to the 1996 Dickey Amendment mandating that the injury prevention funds at Centers for Disease Control and Prevention (CDC) cannot be used to advocate or promote gun control [25]. A 2018 spending bill clarified the language of the Dickey Amendment, allowing the CDC to research gun violence, which was believed to be barred by the amendment, while stipulating that government funds may not be used for gun control advocacy [26]. Additionally, the United States government has no definition for a mass shooting but does define a mass killing as an incident in which a single perpetrator kills at least three people in a public space; this definition is consequently extended to the definition of a mass shooting by various entities that compile data for the purpose of studying mass shootings.

Several private institutes and organizations have established publicly available data repositories that will be used in this study. Four data sets of mass shootings in the United States were utilized, and only events occurring in the continental United States were considered in analyses. The data sets differ in observation periods in addition to their definitions as to what constitutes a *mass shooting*. Data compilation differences lead to large differences in total number of observations. Further discrepancies are acknowledged below and summarized in Table 1. Plots containing the number of mass shootings per month are displayed in Fig 1, with each plot displaying the same time window. Fig 2 displays the distribution of the number of victims from events in each data set.

Table 1. Summaries for each data set used in the analysis. Summaries include the time window of data used in the analyses, definition of what constitutes a mass shooting, and the number of observations falling in the time window.

Data Set	Beginning Date	End Date	Definition	Observations
Brady	February 2005	January 2013	3+ killed	477
Stanford	January 1999	June 2016	3+ shot	262
Mother Jones	January 1999	February 2020	3+ killed	92
GVA	January 2012	August 2020	4+ shot	2669

<https://doi.org/10.1371/journal.pone.0248437.t001>

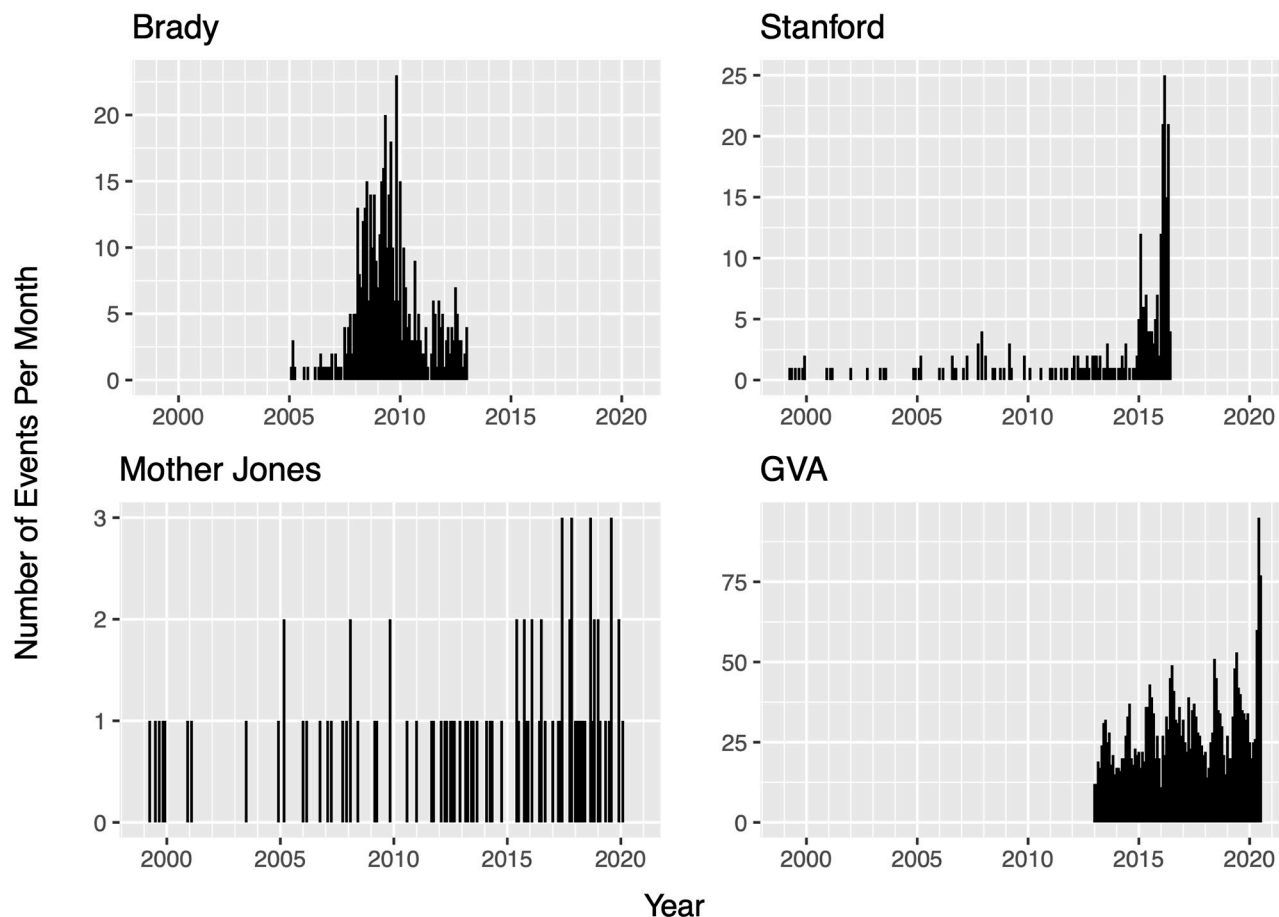


Fig 1. Monthly totals for the number of mass shootings.

<https://doi.org/10.1371/journal.pone.0248437.g001>

Brady: United against gun violence

The Brady Campaign (<https://www.bradyunited.org>) is a nonprofit group advocating for gun control and striving to end gun violence. The organization is named after James Brady, a cabinet member during the Ronald Reagan presidency who was shot during the assassination attempt on the president. Brady, left permanently disabled from the gunshot, became an advocate for gun control. The group has compiled data including incidents in which at least three people were shot or injured, but not necessarily killed. The data spans from February 2005 to January 2013, containing a total of 477 incidents. The Brady Campaign data set used in this article is also used in the Towers, et al.'s [7] analysis to allow for comparison of results. Data can be accessed here: <https://journals.plos.org/plosone/article/file?type=supplementary&id=info:doi/10.1371/journal.pone.0117259.s002>.

Stanford mass shootings in America

The Stanford Mass Shootings in America (<https://library.stanford.edu/projects/mass-shootings-america>) data was compiled in an effort to create a comprehensive collection of mass shooting data in the United States. Incidents included involve three or more people shot, but not necessarily killed. The data ranges from August 1966 to June 2016 when maintenance and updates to the database were halted. We utilize data beginning in January 1999, with

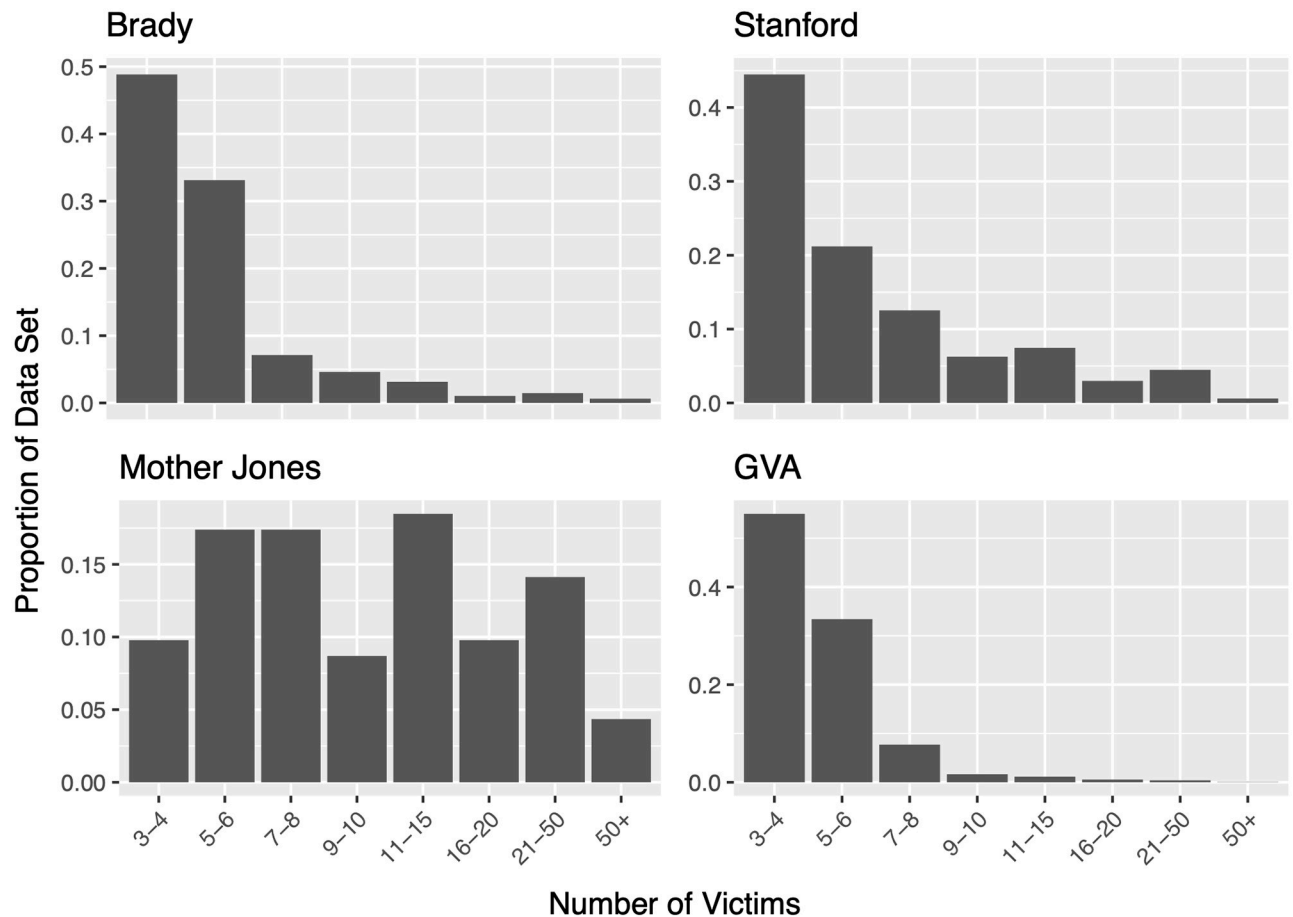


Fig 2. Distribution of the number of victims in each data set.

<https://doi.org/10.1371/journal.pone.0248437.g002>

Columbine happening months later on April 20, 1999, to study the occurrence of mass shootings as a more modern phenomenon. The data originally contained 335 observations, but was reduced to 262 to reflect the altered starting date.

Gun Violence Archive

Gun Violence Archive (GVA) (<https://www.gunviolencearchive.org>) is a nonprofit group that compiles records of gun related incidents in the United States. Incidents recorded involved four or more people shot but not necessarily killed. New records are updated in near real time, with data ranging from January 2012 until the present. While some data sets exclude events such as gang violence, GVA does not set any limiting terms to their definition of a mass shooting other than the number of individuals shot and killed, leading to a data set that contains a greater number of events. For events in which the perpetrator is killed or commits suicide during the shooting, GVA also differs from the other data sets in that the perpetrator is included in the number of total victims.

Mother Jones

Mother Jones (<https://www.motherjones.com/politics/2012/12/mass-shootings-mother-jones-full-data/>) is an investigative journalism organization that has compiled a collection of mass

shootings under stricter criteria than others. With data ranging from 1982 until the present, Mother Jones initially recorded only incidents in which four or more people were killed. When the United States government redefined a mass killing to involve three or more people, Mother Jones followed suit, redefining the criterion for the database. For this analysis, the data will be reduced to events taking place on or after January 1, 1999.

Results

Nonparametric Hawkes processes were fit to each data set using the MISD algorithm to estimate their conditional intensity functions. Initially, the spatial triggering component was included in the conditional intensity function but was later dropped as the spatial triggering component was found to have little to no effect in triggering subsequent events. The remaining results focus on the triggering of the temporal and mark components, $g(t)$ and $k(m)$ respectively. Intervals for the temporal triggering function were chosen to reflect natural breaks in the inter-event time differences, i.e. 2 weeks, 3 months, 6 months, +1 year, while the intervals for the marks triggering function were selected using quantiles to roughly allocate an equal number of events into each interval based on the number of victims. With discrete mark values, an exactly uniform division of events into quantiles could not be realized as certain values accounted for a large proportion of the data that would otherwise have spanned several quantiles, specifically in the GVA data in which roughly 55% of incidents involved four victims.

The diagonal mass of the probability matrix P , estimated background rate, average number of offspring events, and average number of offspring events occurring in the first two weeks are displayed in Table 2. For most data sets, the majority of events are probabilistically treated as triggered events, with background events making up roughly 10% to 54% of observed mass shootings. The estimated background rate for the Brady, Stanford, and Mother Jones data sets are estimated to be between 0.007 to 0.016 mass shooting events per day while the background rate for GVA is substantially larger with an estimated daily rate of mass shootings of 0.33.

For the Brady data set, the model estimated the expected number of offspring per mass shooting event to be roughly 0.90 events with 0.20 of those events, occurring in the first two weeks. This then implies that for an event in the Brady data, 23% of the offspring events occur in the first two weeks with the remaining 77% of events occurring sometime afterward. The Stanford data had an estimated expected number of offspring per event of 0.72 with 60%, 0.43, of these events occurring in the first two weeks. The GVA data set had a slightly smaller overall expected number of offspring per event than Brady or Stanford with an expected number of 0.66 child events. However, the overwhelming majority, approximately 90%, of the offspring events occurred in the first two weeks. Meanwhile, the Mother Jones data had the smallest expected number of offspring per events, 0.46 child events per mass shooting, yet 98% of the child events occurred more than two weeks after the initial mass shooting.

Table 2. Numeric summaries of implementing the MISD algorithm for each data set. Diagonal mass indicates the percent of the probability matrix P that lies on the main diagonal. Background rate is the estimated background rate of the data catalog. Number of offspring is the estimated number of events that are triggered offspring of previous events. Number of 14 day offspring is the estimated number of offspring occurring within 14 days of an event and percent of offspring in first 2 weeks is the estimated percent of an event's total offspring which occurs in the first 14 days after the event.

Data Set	Diagonal Mass	Background Rate	Number of Offspring	Number of 14 Day Offspring	Percent of Offspring in First 2 Weeks
Brady	10.12%	0.0164	0.9007	0.2034	22.58%
Stanford	28.74%	0.0118	0.7153	0.4308	60.22%
GVA	34.14%	0.3272	0.6588	0.5943	90.20%
Mother Jones	54.63%	0.0065	0.4587	0.0111	2.42%

<https://doi.org/10.1371/journal.pone.0248437.t002>

The estimated histogram estimators for the triggering functions of each data set are shown in Fig 3. For each plot, the estimated constants of the histogram estimator step functions are shown as a horizontal line spanning the time or mark sub-interval for which the constant was estimated. The grey vertical bars then represent ± 2 standard errors for each estimated constant of the histogram estimator. The standard error bars are truncated at zero to reflect only values that plausibly represent the phenomenon of interest. The temporal triggering functions, $g(t)$, are densities and thus the areas underneath the step function represent the probabilities of child events occurring over some time-span. The marks triggering functions, $k(m)$, represent productivity multipliers which increase or decrease the rate of triggered events based on the number of victims impacted in prior mass shootings. The x -axes of the temporal triggering functions are truncated as the functions tended towards zero as $t_j - t_i$, for $j > i$, grew larger; x -axes of the marks triggering functions are truncated shortly after the final sub-interval as shown graphically.

In general, with the exception of Mother Jones, the value of the temporal triggering function, $g(t)$, monotonically decreases as t increases to each subsequent time bin. For the Brady data, the decrease in the temporal triggering decreases more smoothly from roughly 0.0161 to 0.0079, to 0.0015 down to approximately 0. For the Stanford and GVA data, the decay in the temporal triggering decreases much more drastically; from 0.043 down to 0.005 down to zero for the first three time intervals in the Stanford data and from 0.0642 down to 0.0013 to essentially zero in the first three time intervals in the GVA data. For the Mother Jones data, the temporal estimates of the triggering are more volatile with estimates starting around 0.0017 and 0.0006 for the first and second time interval, rises to around 0.0036 and 0.0033 in the third and fourth intervals, then finally falls to zero. The Mother Jones data is also unique in that the estimated constants of the triggering function are much smaller in value than the other data sets.

For the estimated triggering functions of the marks for the Brady data, $k(m)$ had an estimated productivity of around 0.6877 for the initial interval, and then increased to 1.4355 for five victims, before falling to 0.9744 for 6-8 victims and 0.5883 for nine or more victims. The estimated mark triggering functions for Stanford and GVA contain the same pattern of an initial increase followed by two descending values. Stanford has an estimate of 0.9218 for the initial bin and then jumped to 1.3817 for five victims, before falling to 0.4053 for 6-7 victims and 0.1323 beyond 7 victims. GVA begins with at 0.5753, increasing to 0.8652 for 5 victims, then falls to 0.7248 and 0.1311 for 6-9 and 10+ victims, respectively. For the Mother Jones data, $k(m)$ also followed a less consistent form, with the highest value of 1.2336 in the first bin before falling to 0.1888 for 7-10 victims and 0.0037 for 11-17 victims before rising to 0.288 for 18 or more victims.

Fig 4 shows the observed number of monthly mass shootings for each data source along with the estimated number of monthly shootings based on the models. The estimated values are computed by taking the median conditional intensity for each month and multiplying it by the length of the month. The models appear to fit the data fairly well in that the estimated number of monthly mass shootings tends to follow the trends in the observed values. Both the Mother Jones and Stanford data sets contain instances where no mass shooting events occurred over a sequence of consecutive months. For these months, the models tended to over-estimate the number of events as the model assumes a constant background rate.

Super-thinning was implemented to evaluate each model's fit to the individual data sets with tuning parameter, b , set to the median estimated conditional intensity for each source. To assess the overall fit of the model, the residual process for each data set is displayed as histograms in Fig 5. If the model fits the data well, then we would expect the histograms to demonstrate a roughly uniform distribution throughout the entire time window. Of the four data sets, the estimated model for the Mother Jones data appears the least uniform in shape with

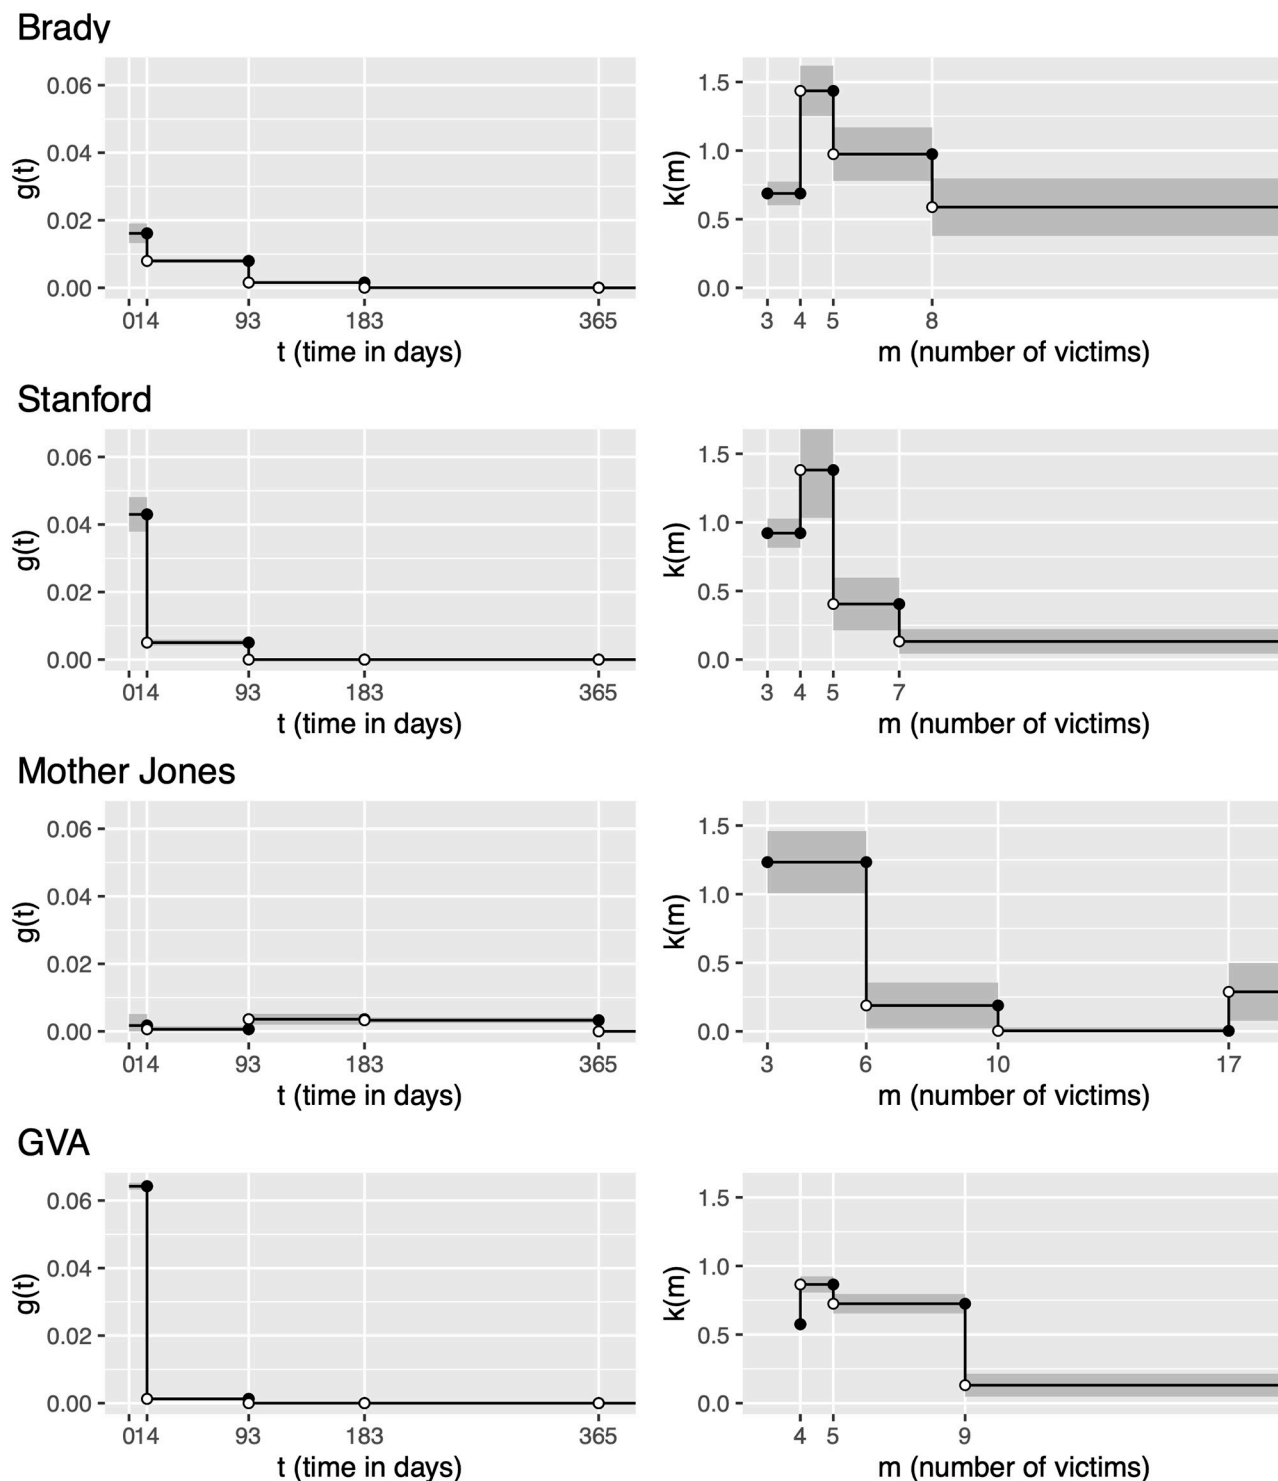


Fig 3. Estimated histogram estimators. Figures on the left display estimated values of the temporal triggering function plotted over time, with the time bins used in the analysis shown on the x axis. Figures on the right display values of the marks triggering function over the set of marks (number of people injured or killed). Standard error regions are shown in gray and x -axes are truncated after the largest intervals.

<https://doi.org/10.1371/journal.pone.0248437.g003>

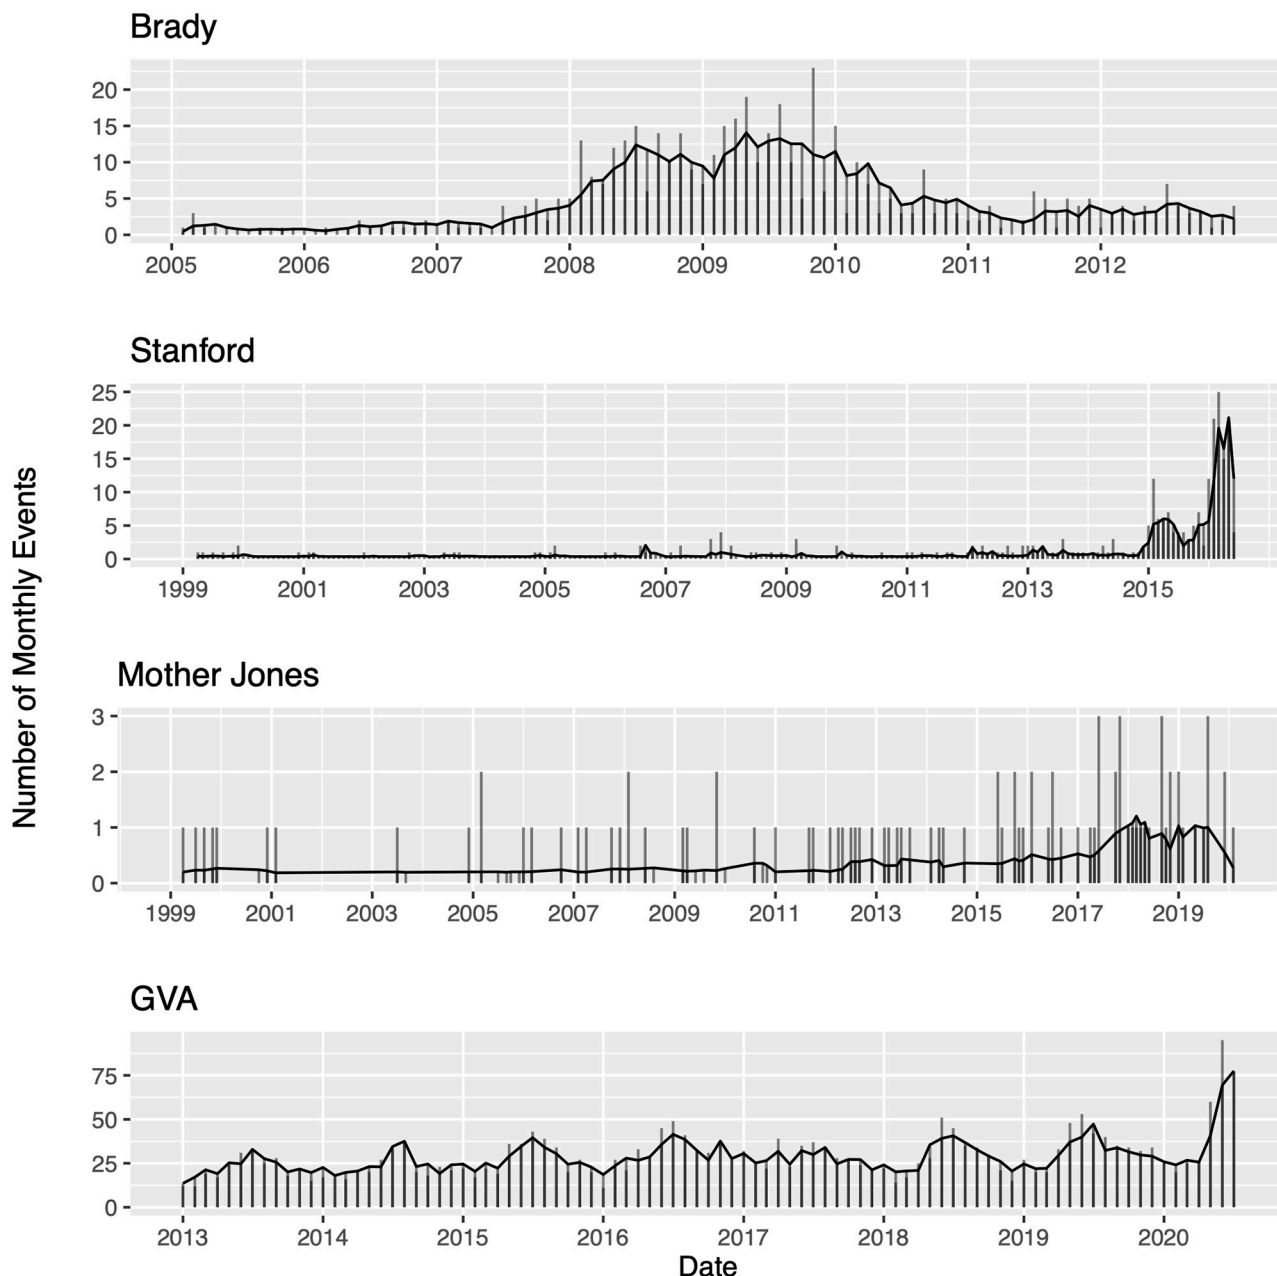


Fig 4. Estimated conditional intensities. The number of monthly mass shootings are plotted as vertical lines over time. The median value of the estimated conditional intensity of the observed points is calculated for each month, multiplied by the number of days in each corresponding month, and plotted (solid line) over time.

<https://doi.org/10.1371/journal.pone.0248437.g004>

substantial deviations throughout the time-window. The residual process for the GVA data appears the most uniform overall, though also with some deviations. The distributions of the Brady and Stanford residual processes are somewhere in the middle with many time intervals appearing roughly uniform with some systematic deviations for certain time periods. The residual process for the Stanford data appears to have, in general, lower values prior to 2005 and slightly higher values in the years following, while the Brady residual process exhibits more of a unimodal distribution with a peak in values from 2008–2010.

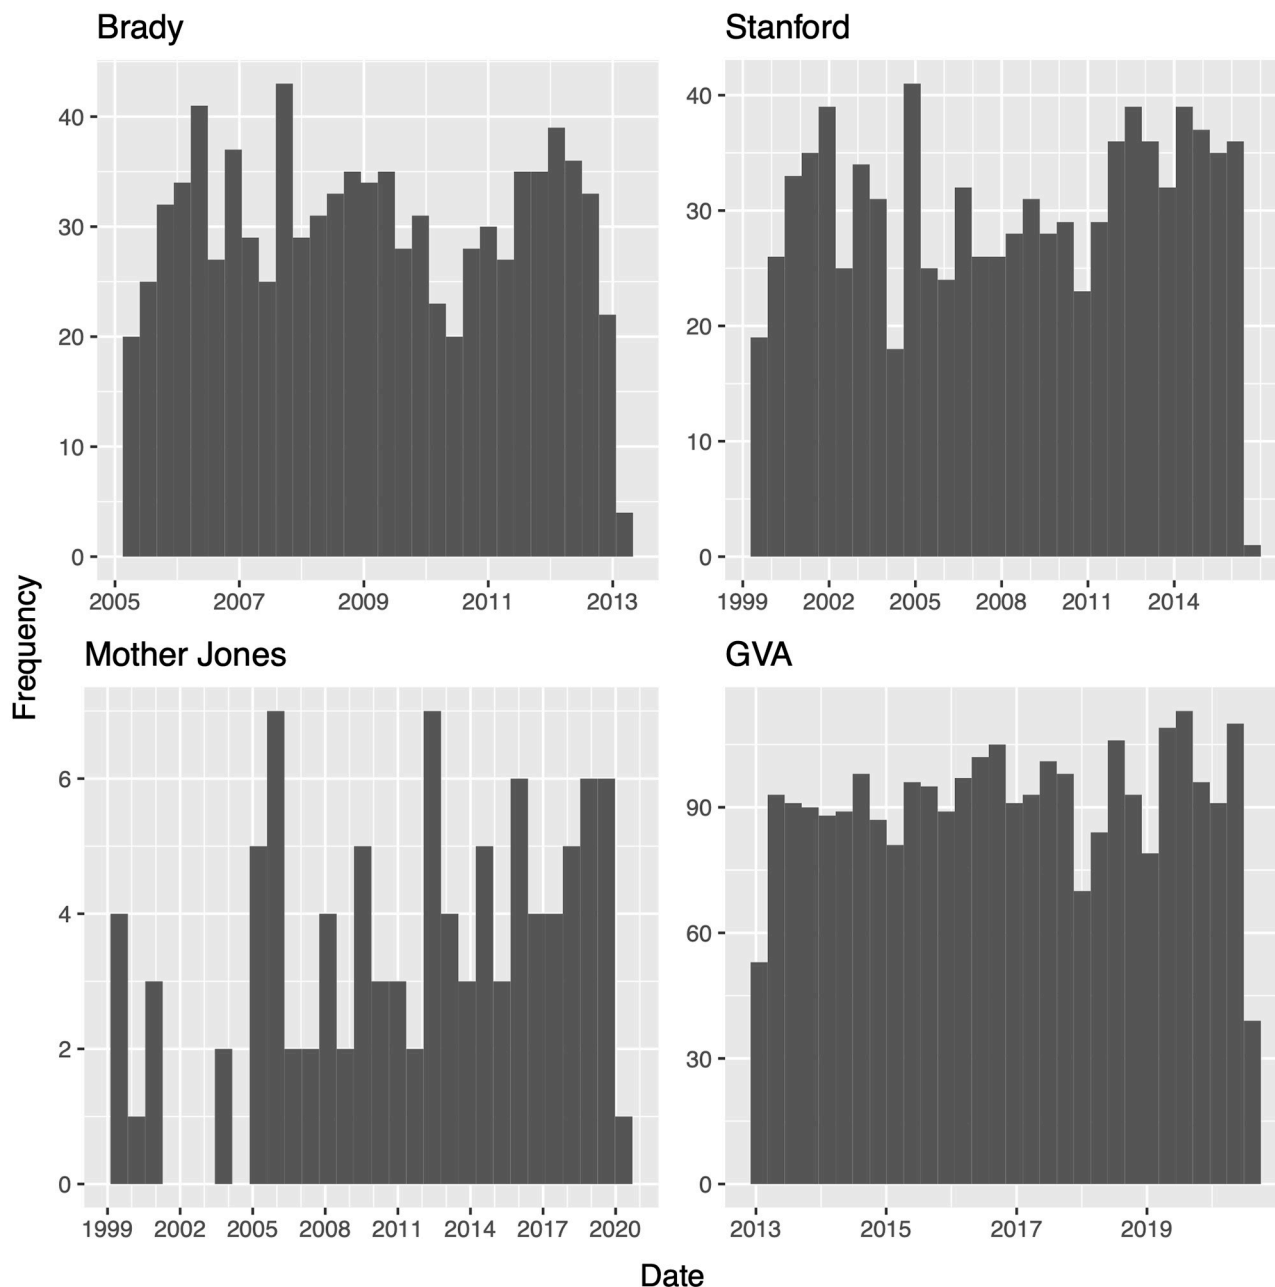


Fig 5. Histograms of super-thinned processes. After super-thinning is implemented, the data are plotted over time, displaying the distribution of the super-thinned process.

<https://doi.org/10.1371/journal.pone.0248437.g005>

Discussion

In this article, we investigate the contagiousness of mass shootings by treating the data as a marked self-exciting point process and analyze it through nonparametric Hawkes procedures. The contagiousness of mass shootings was previously studied by Towers, et al. [7], reporting that each mass shooting will incite at least 0.30 new events brought on by an increase in probability of events that lasts for 13 days, on average, after an event. The self-excitation contagion model utilized in the Towers analysis requires several parametric assumptions including

assuming a distribution for the decay of contagiousness, a constant number of secondary events, and the duration of contagion process. With little research on the contagiousness of mass shootings, circumventing the reliance on parametric assumptions through a nonparametric modeling framework is an important contribution to the study of this devastating phenomenon.

Similar to the results found in the Towers, et al. article [7], we found no substantial spatial effect using the nonparametric framework. We do however find evidence that events may produce higher numbers of offspring than previous results, with estimated number of offspring ranging from 0.45 to 0.89, as much as almost 3 times the value reported by Towers, et al. when using the same Brady Campaign data set. We also note that the contagion effect estimated by the nonparametric model, in time, lasts longer than the 13 days found using the previous parametric model, particularly in the Brady data set, with roughly 77% of triggered events occurring after two weeks.

These differences between the parametric and nonparametric approaches are likely due to the prescribed functional form for the temporal decay of the triggering function in the parametric model. The parameter controlling the rate of exponential decay of the parametric model is able to capture the contagion effect in the first two weeks, where the effect is strongest, but then decays too quickly and does not capture the long-term contagion effect. In other words, the exponential form of decay does not have a fat enough tail to estimate long-term contagion effects while also capturing the near-term effects, leading to an underestimation of the overall amount of contagion. Utilizing the nonparametric model resolves this issue by estimating the length and amount of contagion using only the data without the need to specify a functional form for the decay.

In Fig 3, the temporal histogram estimators tended to agree that the initial two-week period after a mass shooting event tended to have larger contagion effects compared to time periods after the initial two weeks, save for Mother Jones which had a temporal histogram estimator which was much more volatile. This volatility might not be entirely unexpected given that the Mother Jones data set had slightly more than one-third of the total number of observations compared to the next smallest data set despite the catalog spanning the full temporal window used in the analyses. These factors then imply that very few of the pairwise time differences between events in the Mother Jones data fall in the shorter time intervals. The GVA data meanwhile is by far the largest data source with the shortest time-window. For the GVA data, nearly all of the contagion factor occurs in the first two-weeks. This is likely due to many of the pairwise inter-event time differences occurring relatively quickly after previous events.

The triggering functions for the marks show much less consistency between the data sets but demonstrate the benefit of allowing the expected number of secondary events to vary depending on the size of the marks. The histogram estimator for the number of victims for the Brady campaign, demonstrates that mass shootings with larger numbers of victims increases the productivity of those events in spurring future events. The Stanford data meanwhile shows that events with between four to seven victims were more productive than larger events with greater than seven victims, while a similar result was seen in the GVA data. Again, the histogram estimator for the Mother Jones data is drastically different compared to the rest in that smaller events are more productive than larger events with high victim numbers. Furthermore, while the model appears to be finding some signal in regards to how the number of victims impacts the productivity of mass shooting events to spur future events, it should be noted that there's a considerable amount of uncertainty in these estimates, as represented by the standard error bars, especially for the Stanford and GVA data.

Fig 5 shows the results of super-thinning the point process models for the different data catalogs. By considering the uniformity of the super-thinned residual processes we can evaluate

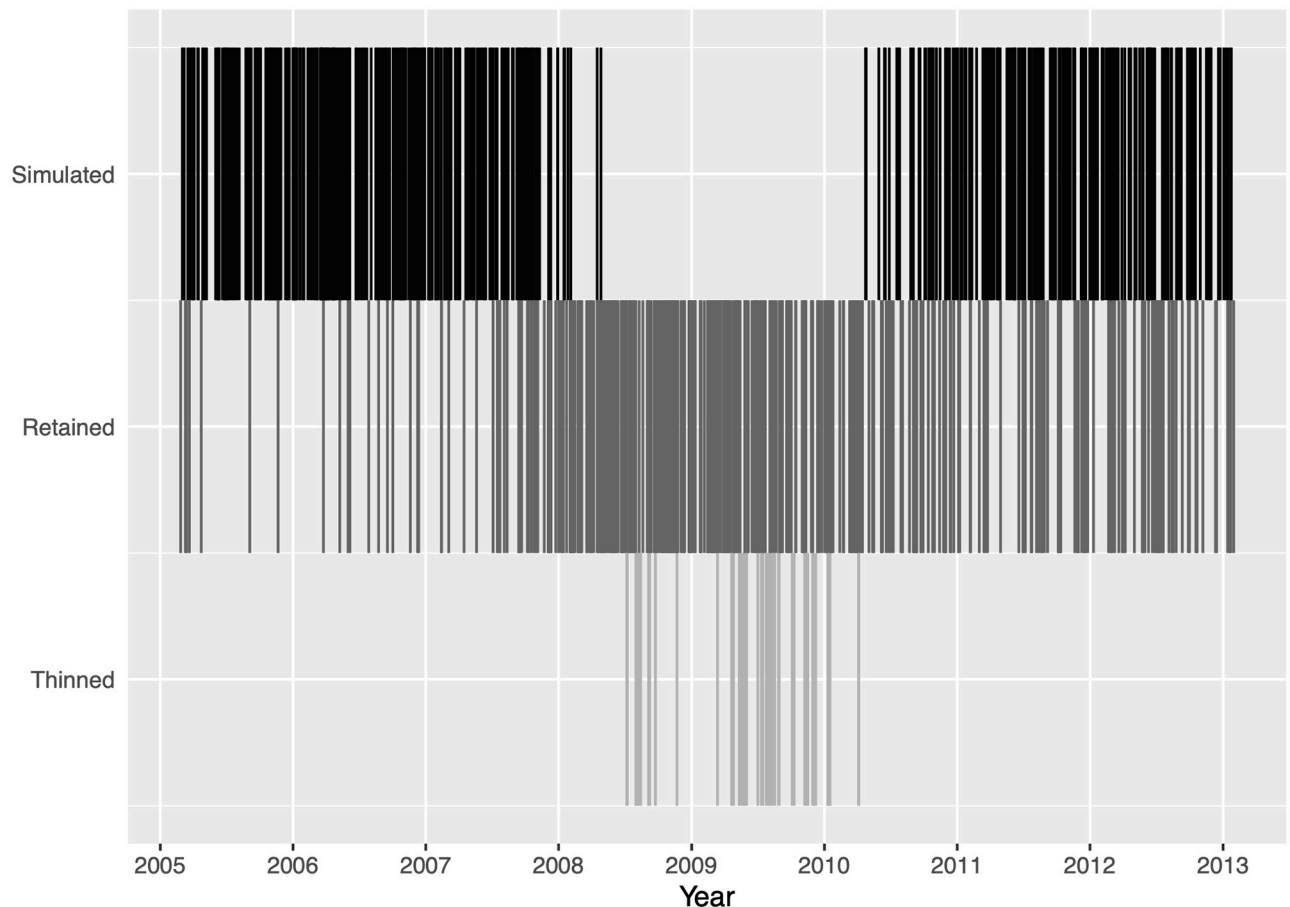


Fig 6. Brady campaign super-thinning plot. After super-thinning is implemented, events are plotted by their classification type over time, indicating events that were removed (thinned) from the data, events that were not removed (retained), and simulated events were superposed into the data (simulated).

<https://doi.org/10.1371/journal.pone.0248437.g006>

the overall fit of the models in that models that fit the data well should have a uniform appearance in the histograms.

We observe that the residual process for the Brady model has a unimodal appearance rather than the desired uniform distribution. Examining Fig 6, which shows the composition of the points for the super-thinned residual process for the Brady data, allows us to further investigate the unimodal distribution. The simulated lines at the top of the plot shows the points which were superposed while the retained lines show the points of the original process which were retained after thinning. The points which were thinned are then shown at the bottom of the plots. From the figure, it is evident that the super-thinned process simulates events in areas of low intensity and removes events from areas of high intensity, but by simultaneously analyzing figures the histogram for the Brady data of Figs 5 and 6, we see that a lack of sufficient thinning spurs departures from uniformity in the histogram. This lack of thinning then indicates that the model was not able to capture the full contagion effect present in the data.

In Fig 5, we observe an approximately uniform distribution for the Stanford residual process, save a few spikes and falls, most notably at the end of 2010. Throughout the Stanford catalog, super-thinning appears to be performing as expected, despite the abrupt increase in the number of shootings that can be seen in the Stanford data's conditional intensity. Fig 5 also displays an approximately uniform distribution after super-thinning the GVA catalog.

Fig 5 shows a non-uniform distribution of super-thinned residuals for the Mother Jones data. With so few events recorded in the Mother Jones data set, well-fitting models are more challenging to realize without adding further complexity to the model. In Fig 4, the frequency of observed events in the Mother Jones catalog appears to vary considerably over time, with 40% of events occurring in only the last five years of the catalog. With such disparity in the frequency of events, fitting a single background rate for the entire process may oversimplify trends in the data; employing a nonconstant background rate may allow for a stronger representation of the data.

Amongst the data sources used for this article, we see in Table 1 that definitions used by the different sources to define *mass shootings* does not vary dramatically, yet Fig 1 reveals stark differences in the timing and number of events in each catalog. Some sources, such as the Gun Violence Archive, were very permissive in attributing the *mass shooting* label to gun crimes. Other sources, such as Mother Jones, were much more stringent. These inconsistencies between data sources lead to numerical results which vary, sometimes substantially, across analyses.

To further highlight the inconsistencies evident in the different data sources, we note that the Gun Violence Archive reports 2,669 mass shootings in less than nine years, while the original Mother Jones data reports 118 incidents over nearly thirty-eight years. And although Brady and Mother Jones both define mass shootings as events in which three or more individuals are killed, the number of events in each data set are starkly different in the overlapping time period of 2005 to 2013. Such issues which might only be remedied by the creation of a more consistent definition and definitive catalog of mass shooting events.

Conclusion

In this article, we assess the contagiousness of mass shootings using a nonparametric Hawkes process framework for a variety of data sources. This framework relies on fewer parametric assumptions than previous studies and detects a contagion effect which varies over both time and the number of victims. We also find that the level of contagion is contingent upon the data source used as no definitive catalog of data for mass shootings yet exists.

Although the estimated conditional intensity for each process appears to closely mirror the true data process, more complex models with additional features may yield better fitting models in the future. Specifically, adapting a nonconstant background rate over time and/or a productivity function which is allowed to vary over time would allow future models to capture temporal changes to these two components. More complex models might also allow for the incorporation of meaningful spatial attributes or additional relevant covariates. The models featured in this article then represent a baseline approach for the modeling of mass shootings as the nonparametric framework we implemented is extensible and able to benefit from innovations made in other fields and applications.

Supporting information

S1 Data.
(ZIP)

Acknowledgments

The authors would like to thank the following organizations for making their data available:

- Brady Campaign to Prevent Gun Violence, <https://www.bradyunited.org/>

- Stanford Mass Shootings in America, courtesy of the Stanford Geospatial Center and Stanford Libraries, <https://library.stanford.edu/projects/mass-shootings-america>
- Mother Jones, <https://www.motherjones.com/politics/2012/12/mass-shootings-mother-jones-full-data/>
- Gun Violence Archive, <https://www.gunviolencearchive.org/>

The authors would also like to thank Patrick Cummings for his guidance and suggestions in implementing the estimation algorithms.

Author Contributions

Conceptualization: James Molyneux.

Data curation: Peter Boyd, James Molyneux.

Formal analysis: Peter Boyd, James Molyneux.

Investigation: Peter Boyd.

Methodology: Peter Boyd, James Molyneux.

Software: Peter Boyd, James Molyneux.

Supervision: James Molyneux.

Visualization: Peter Boyd.

Writing – original draft: Peter Boyd, James Molyneux.

References

1. Bauchner H, Rivara FP, Bonow RO, Bressler NM, Disis MLN, Heckers S, et al. Death by Gun Violence—A Public Health Crisis. *JAMA Psychiatry*. 2017; 74(12):1195–1196. <https://doi.org/10.1001/jamapsychiatry.2017.3616> PMID: 29052736
2. Grinshteyn E, Hemenway D. Homicide, Suicide, and Unintentional Firearm Fatality: Comparing the United States With Other High-Income Countries, 2003. *The Journal of trauma*. 2011; 70:238–43. <https://doi.org/10.1097/TA.0b013e3181dbadbf>
3. Meindl J, Ivy J. Mass shootings: The role of the media in promoting generalized imitation. *American Journal of Public Health*. 2017; 107(3):368–370. <https://doi.org/10.2105/AJPH.2016.303611> PMID: 28103074
4. Dzau VJ, Leshner AI. Public Health Research on Gun Violence: Long Overdue. *Annals of Internal Medicine*. 2018; 168(12):876–877. <https://doi.org/10.7326/M18-0579> PMID: 29554693
5. Federal Bureau of Investigation. 2000 to 2018 Active Shooter Incidents; 2016. Available from: <https://www.fbi.gov/file-repository/active-shooter-incidents-2000-2018.pdf/view>.
6. Jetter M, Walker JK. The Effect of Media Coverage on Mass Shootings. Institute of Labor Economics (IZA); 2018. 11900. Available from: <https://ideas.repec.org/p/iza/izadps/dp11900.html>.
7. Towers S, Gomez-Lievano A, Khan M, Mubayi A, Castillo-Chavez C. Contagion in Mass Killings and School Shootings. *PLOS ONE*. 2015; 10(7):1–12. <https://doi.org/10.1371/journal.pone.0117259> PMID: 26135941
8. Marsan D, Lengliné O. Extending Earthquakes' Reach Through Cascading. *Science*. 2008; 319(5866):1076–1079. <https://doi.org/10.1126/science.1148783> PMID: 18292339
9. Daley DJ, Vere-Jones D. An Introduction to the Theory of Point Processes Volume I: Elementary Theory and Methods. Springer Science & Business Media; 2004.
10. Hawkes AG. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*. 1971; 58(1):83–90. <https://doi.org/10.1093/biomet/58.1.83>
11. Ogata Y. Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical association*. 1988; 83(401):9–27. <https://doi.org/10.1080/01621459.1988.10478560>

12. Ogata Y. Space-time point-process models for earthquake occurrences. *Annals of the Institute of Statistical Mathematics*. 1998; 50(2):379–402. <https://doi.org/10.1023/A:1003403601725>
13. Fox EW, Schoenberg FP, Gordon JS. Spatially inhomogeneous background rate estimators and uncertainty quantification for nonparametric Hawkes point process models of earthquake occurrences. *Ann Appl Stat*. 2016; 10(3):1725–1756. <https://doi.org/10.1214/16-AOAS957>
14. Zhao Q, Erdogdu MA, He HY, Rajaraman A, Leskovec J. SEISMIC: A Self-Exciting Point Process Model for Predicting Tweet Popularity. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD'15*. New York, NY, USA: Association for Computing Machinery; 2015. p. 1513–1522. Available from: <https://doi.org/10.1145/2783258.2783401>.
15. Mohler GO, Short MB, Brantingham PJ, Schoenberg FP, Tita GE. Self-Exciting Point Process Modeling of Crime. *Journal of the American Statistical Association*. 2011; 106(493):100–108. <https://doi.org/10.1198/jasa.2011.ap09546>
16. Porter MD, White G. Self-exciting hurdle models for terrorist activity. *Ann Appl Stat*. 2012; 6(1):106–124.
17. Gerhard F, Deger M, Truccolo W. On the stability and dynamics of stochastic spiking neuron models: Nonlinear Hawkes process and point process GLMs. *PLOS Computational Biology*. 2017; 13(2):1–31. <https://doi.org/10.1371/journal.pcbi.1005390> PMID: 28234899
18. Meyer S, Elias J, Hohle M. A Space-Time Conditional Intensity Model for Invasive Meningococcal Disease Occurrence. *Biometrics*. 2012; 68(2):607–616. <https://doi.org/10.1111/j.1541-0420.2011.01684.x> PMID: 21981412
19. Omori F. On the Aftershocks of Earthquakes. *Journal of the College of Science, Imperial University of Tokyo*. 1894; 7:111–120.
20. Utsu T. A statistical study on the occurrence of aftershocks. *Geophys Mag*. 1961; 30:521–605.
21. Gutenberg B, Richter CF. Frequency of earthquakes in California. *Bulletin of the Seismological Society of America*. 1944; 34(4):185–188.
22. Clements RA, Schoenberg FP, Veen A. Evaluation of space–time point process models using super-thinning. *Environmetrics*. 2012; 23(7):606–616. <https://doi.org/10.1002/env.2168>
23. Schoenberg FP. Multidimensional Residual Analysis of Point Process Models for Earthquake Occurrences. *Journal of the American Statistical Association*. 2003; 98(464):789–795. <https://doi.org/10.1198/016214503000000710>
24. Brémaud P. *Point Processes and Queues*. Springer series in statistics. Springer New York; 1981.
25. Rostron A. The Dickey Amendment on Federal Funding for Research on Gun Violence: A Legal Dissection. *American journal of public health*. 2018; 108(7):865–867. <https://doi.org/10.2105/AJPH.2018.304450> PMID: 29874513
26. DeBonis M, O'Keefe E. Here's what Congress is stuffing into its \$1.3 trillion spending bill. *The Washington Post*. 2018;.